

Bridging belief function theory to modern machine learning

Thomas Burger*

iRTSV-BGE (Université Grenoble-Alpes, CNRS, CEA, INSERM)

Grenoble, France

thomas.burger@cea.fr

April 16, 2015

Abstract

Machine learning is a quickly evolving field which now looks really different from what it was 15 years ago, when classification and clustering were major issues. This document proposes several trends to explore the new questions of modern machine learning, with the strong afterthought that the belief function framework has a major role to play.

Keywords: belief functions; machine learning;

1 Introduction

In an age of user generated web-contents and of portable devices with embedded computer vision capabilities, machine learning (ML) and big data mining questions are fundamental. As a result, these questions naturally penetrate neighboring research fields, including belief function theory (BFT), so that it is now usual to attend a “Classification” session [26] or a “Machine Learning” session [16] in a conference devoted to belief functions.

However, it is hard to accept that among the various proposed approaches based on BF, very few have become state-of-the-art ML methods, the knowledge of which has spread beyond the BF community. Without any doubt, this can be partly explained by the relative size of the scientific communities under consideration: although quickly growing, the BF one is relatively small with respect to that of statistics, Bayesian networks, neural networks, etc. However, this reason alone is not sufficient: There are indeed other topics, such as for instance, information fusion, where BF-based methods are now as well recognized as are methods based on more classical formalisms, such as probabilities, or ontologies.

In this report, I assume an additional reason: that some researchers focused on BFT (especially the youngest), who have progressively turned their interests towards ML problems, may not capture the newest trends of this field. In fact, I used to be an example of such researchers, and I acknowledge that my first perceptions of ML were clearly outdated. This is why, I propose a short review of the respective evolution of BFT and of ML, as well as an attempt to put them in perspective. Of course, many senior researchers may find this exercise futile, as they have their own broad view on the question. However, to my knowledge, no recent referenced article is available for any reader seeking for a starting point to question the links between ML and BFT.

This document is structured as follow: In Section 2, a brief recall of the evolution of the mainstream in the BF community is provided. Then, in Section 3, a short summary of the earlier ages of ML up to the mid-90s, is sketched, as well as a coarse description of the successful interactions between ML and BFT in those times. Afterward, I provide in Section 4 a synthetic overview of the revolution that blew over ML around the early 2000s, and which modified its goals and the organization of its supporting community. As BFT does not seem to fit in this new picture of the ML world, I list in Section 5 a few problems that may still be of interest for the current mainstream of BFT, as well as some potential interesting evolutions for the community to adapt to the newly raised questions.

2 The landscape of belief function interpretations

When considering BF interpretations, one often opposes Dempster’s *imprecise statistic view* to Shafer and Smets *singular view*, such as described in [27]. However, there is another way to sort the various interpretations: it is to refer to the mathematical object a mass function is perceived as a generalization of. This leads to the following taxonomy:

*The author would like to thank Fabio Cuzzolin, who organized, a “*discussion panel on the status and future of the belief function theory*” during the BELIEF 2014 conference [16]. This very enriching brainstorm was the starting point for this work, which can be seen as a rejoinder to this discussion panel.

- The **probability-affiliated** interpretation: “A mass function is an object which generalizes a discrete probability measure, where the probability masses are not necessarily known, but are only assumed to belong to an interval”. Whatever the origin of this imprecise probability, i.e. either statistical (such as in Dempster’s view [19, 20], or as in the Theory of Hints [39]), or subjective (such as with Shafer or Shenoy [55, 58, 59, 57]), the mathematical theory behind is that of random sets [36, 43].
- The **set-affiliated** interpretation: “A mass function is a set description of the real value of an ill-known variable, which is enriched by weightings which add up to one”. Of course, this view is conceptually closer to that of fuzzy sets, possibility theory [31], and to artificial intelligence (AI) in general.
- The **capacity-affiliated** interpretation: “A mass function is only a particular type of Choquet’s capacity, the Moebius transform of which is totally monotone”. Such interpretation is strongly related to multi-criteria optimization and to operational research [33].

From a historical perspective, the oldest probability-affiliated interpretation (of Dempster) was rooted in imprecise statistics, before moving towards non-frequentist views (that are more classically associated to nowadays subjective views), under the influence of Shafer; yet, in Shafer and Shenoy works [56], the link to probability remains strong as the reference to random set is explicit. Later on, the set-affiliated interpretation has spread under the influence of renown authors (such as Zadeh, Dubois, Smets or Yager), with a strong background in data fusion, fuzzy sets and expert systems. Among them, Smets constantly positioned his work with respect to Bayesian classifiers [51] or Bayesian reasoning [63]; however, the Transferable Belief Model (TBM [60]) was a non-probabilistic model, which, in Smet’s view [61], was not compliant with random sets. Besides, the capacity-affiliated interpretation of BF has never really developed on its own, independently of other works in operational research and in game theory. Finally, todays, the great majority of the belief function community accepted the set-affiliated interpretation, as it appears in proceedings such as [26, 16], where the majority of the articles undertook this interpretation.

3 Classification, clustering and belief functions

As fully described by Miclet and Cornuéjols in [14], the short history of ML is full of rapid and strong evolutions, so that to date, the discipline which was originally part of AI, has completely separated from its parent discipline to make its own path, which appears to strongly converge toward statistics and optimization: They review the original heuristics motivated by bio-inspiration (such as multi-layer perceptrons for instance), their transformation into more systematic explorations (through the concept of version space, introduced by Mitchel [42] in 1982), the golden age of symbolic learning, and finally, the opposition between supervised and unsupervised learning, so that the classification and clustering problems were at the center of attention in the mid-90s.

Yet it is not described in [14], Probabilistic Graphical Models (PGM) developed in parallel during this period. In fact, this omission is not really surprising, as PGM is a community on its own, well separated from the ML community. In fact, PGM are also classically used in problems such that data fusion or system diagnosis, which may not involve any learning. However, for a long time, Bayesian networks were amongst the state-of-the-art supervised learning methods, along with others non-probabilistic methods (k -NN, decision tree, etc.). On the other hand, unsupervised learning was mainly non-probabilistic: AHC, Kohonen maps, and of course, k -means. This latter is interesting: Although not probabilistic, this algorithm provides a relaxed solution of the EM algorithm applied to a mixture of Gaussian models [21, 1]. However, in parallel, Bezdek proposed [5], and established [44] a fuzzy version of the k -means, named “fuzzy c -means”, which in spite of its non-probabilistic motivation, behaves similarly to the EM algorithm. This clearly illustrates that, in spite of different culture (statistics or fuzzy sets), similar tools are proposed in different communities. Another episode related to the multiple cultural anchors of k -means recently showed up: it was established that k -means is the discrete counterpart to the PCA [30], a classical non-probabilistic method in statistics, which was developed in the multivariate analysis school [35]. In a nutshell, up to the mid-90s, the connections between ML and IA, although weakening, have remained apparent; The problems of interest in ML are related to clustering and classification; Various formalisms are involved in ML, including probabilities, and none of them claims a clear superiority.

In this context, BFT (still rooted in the probability-affiliated interpretation, while the set-affiliated one is growing) naturally join the trend, so that numerous evidential versions of classical algorithms are successfully proposed and rapidly become state-of-the-art references. Among them, a significant proportion are proposed by Denoeux, such as the evidential k -NN [22] and the evidential neural network [23] in the late 90s or the early 2000s. On this basis, during the early 2000s, numerous derivations of virtually any classical ML algorithms are proposed. Among them, Smets’ series on target identification and on Kalman filter [51, 50, 49, 52, 46, 64] is of prime importance for two reasons. First, they had thrived on top of the generalized Bayesian theorem, proposed a decade earlier [62]. This theorem proposes to link the likelihood and plausibility functions, so that it becomes possible to derive algorithms in the BF framework, which fit the parameters of a model to observed data, in a perfectly compatible ML language [68]. This path was taken over by Denoeux in order to extend Dempster’s EM algorithm [13, 24, 25], while cautiously stepping out of the TBM framework.

The second reason of the importance of Smet’s series on Kalman filter is cultural: As a defender of his own TBM, he constantly opposed the TBM and the Bayesian views [51, 63], so that, under his major influence, ML, which from the mid-90s, is less and less influenced by probabilities, was depicted in the BF community

mainly as a Bayesian field. In this context, numerous researchers of the BF community, including me (as a PhD student of the mid-2000s), were frozen in a decade-old past, where the ML community would be both inspired by Bayesianism and mainly focused on classification and clustering; two facts that obviously do not hold anymore.

4 A recap of modern ML

The works of Vapnik, on support vector machines [7] and on statistical learning theory [69], have provided the foundations of a major revolution which transformed ML in the late 1990s, and which, according to Stéphane Mallat [41], still had major impacts fifteen years later. The cornerstones of this revolution are three of them. The first one is obviously the kernel trick [53], which deeply roots machine learning in the frameworks of distance geometry, Riemannian analysis, and multivariate analysis (a field of statistics which does not assume any probabilistic model underlying the data). The second one is the idea that a learning problem should be addressed through the minimization of its empirical risk [2]. Finally, the third one is to accept that, depending on the dimensionality of the description space, and on the size of the dataset, the empirical risk minimization may be ill-posed, so that a regularizer should be involved in the optimization.

These ideas percolated in the ML community for a decade, providing the tools to give a unifying description [41, 12, 70] of wavelet transform based signal processing, diffusion process, kernel machines and deep neural networks¹. Combined with the first works on variable section by ℓ^1 penalty [67], this lead to major breakthroughs in sparse learning [40], which is of prime importance for the uncertainty theory communities, for its main connection to the *Robust Uncertainty Principles*, proposed by Candès, Romberg and Tao [10].

Todays, in the mid-2010s, ML is not anymore an inter-disciplinary field on which interfere different theories ranging from cognitive sciences through AI to probability, to solve supervised or unsupervised problems: both the background culture and the objectives has changed. Regarding the objectives, they relate to those of information retrieval, social networks, recommender systems, feature extraction, variable selection, data factorization, and sublinear optimization. Even if improving clustering and classification method still deserves some interest², the harder problems presented in ML challenges [3] receives most of the focus. Regarding the background culture, the field is less interdisciplinary and it is mainly considered as a sub-domain of applied mathematics build on top of optimization (mainly convex), geometry, multivariate statistics and harmonic analysis. Very little room is left for subjective probabilities or for AI. This is described in [14], yet it is also well illustrated by the applied mathematics background of most of the researchers recently hired in ML labs.

5 Some room for belief functions in ML

In this context, it may appear as particularly difficult for the BF community to adapt to the recent evolution of ML, in order to provide state-of-the-art developments. First, BF interpretations and ML have had opposite evolutions so far: From a probability-affiliated view compatible with statistics, BFT had given more room to subjectivism, to finally end in a preponderant set-affiliated view tailored to data fusion and AI problems. On the other hand, ML has quit IA to become more and more tied to functional analysis and convex optimization. This antagonism is well illustrated by the following observation: In ML the entries are raw observations, modeled by a set of points living in a vector space; On the other hand, in the TBM, the entries are assumed to be subjective opinions from different agents and of high semantics. In a similar way, a major asset of the BFT is to provide a rich description of the various types of uncertainty associated with some pieces of information; On the other hand, in ML one is classically not interested in modeling them, but rather to blindly minimize a loss function. Nevertheless, BF community still has several cards to play with respect to ML. In the sequel of this section, I present some of them, sorted according to the BF interpretation they affiliate to.

5.1 Staying on to set-affiliated interpretations

The first one is to keep the set-affiliated interpretations, including the TBM, and to restrict to some very specific ML problems where it is adapted. This is definitely the easiest way, as it prevents any change of the current mainstream of the BF community. However, it remains of interest, even if the path is narrowing under the pressure of big data constraints. Obviously, one must focus on problems where, at the first stage, several agents are used on dedicated learning tasks, and at the second stage, some cooperation or combination between them is expected. This encompasses a wide class of problems that are classically faced in computer vision, among which a few of them [17, 18, 48] are addressed in the BFT:

- **Ensemble learning:** the idea is here to combine the capabilities of several classifiers so that their consensus decision is more robust. This setting has long been explored in the BF framework (see [38, 47] as well as their references), by several researchers, including myself [9, 37]. However, the impact all these works is limited by lack of theoretical performance guaranties, such as with boosting methods [32].

¹The latters have revolutionized computer vision as well as many other ML application fields, as described by De Freitas in a recent keynote lecture [16].

²More precisely, clustering and classification challenging problems still exist, yet, in a setting that differs from the original one on which BFT is classically used: for instance classification of billions of items over millions of classes, or computer vision problems where the classification itself is not the issue with respect to the feature extraction problem that precedes.

- **Co-training:** In this setting [6], various classifiers work on different feature spaces or on different datasets, in order to have complementary knowledge. Then, each classifier is used to label examples that are useful for the other classifiers to improve their performances. This can be extended to **transfer learning** problems, where models trained on a first setting are transposed to other similar settings [15].
- **Active learning:** A classifier asks a human agent to label the training examples that it knows are useful to improve its prediction capability [54].

In all these settings, it is required to have several agents, either humans or machines, to automatically evaluate the level of knowledge of each, and to provide a communication scheme between them, so that they can improve one another both their specificity (i.e. being capable of taking a decision), and their consistency (in order to limit their misleading predictions). Described in such a way, BFT seems to be an adapted framework to consider this problem in the most general case. Even if strictly speaking, the learning process would not be accounted for, such a “cooperation framework” would be of prime interest for numerous tasks such as complex scene analysis, bioinformatics, etc.

5.2 Going back to probability-affiliated interpretations

The second option relies on reversing on purpose the evolution of the BF theories, and to go back to the probability-affiliated interpretations, as they root on solid statistical foundations. Even if ML is more and more involved in optimization, the problem formulation remains in the language of statistics, and as such it is compliant with Dempster’s original view³. In such a setting, it could be interesting to address several questions. For instance:

- Is it possible to reformulate the empirical risk minimization principle in the BF framework? Does the minimizer remain convex, if several forms of uncertainties are distinctly accounted for in its expression?
- Can we relate the bias-variance dilemma to that of the specificity-consistency trade-off [45]?
- Is it possible to extend the multiple test correction setting at the center of most omics studies [4, 65] to BF, to account for badly imputed data?

Apart from these rather general questions, let us note that the recent works [13, 24, 25] which rely on the likelihood interpretation of plausibility, in a setting which differs from the TBM is already a step in this direction. The citation rates of these works acknowledge the idea that going back to a probability-affiliated interpretation makes sense.

5.3 Moving on to capacity-affiliated or other interpretations

The last solution is to try to accelerate the evolution of BF community, so that after the probability-affiliated and the set-affiliated interpretations, that of capacity-affiliated focuses the interest. In fact, this view is much more related to optimization, and the options it provides in terms of modeling as well as in terms of solver could definitely be of interest for ML. Basically, it would lead to consider ML learning problems which rely on a game theory setting, or in which multiple criteria have to be optimized in the meantime. Of course, this would naturally lead to models which do not restrict to totally monotone capacities (i.e. BF), but to other types of capacities. More generally, encompassing the BFT in a wider frameworks (Choquet capacities, Walley’s lower capacity, or any other, whatever it is) is potentially enriching, and this line has already been adopted by other researchers on questions which are central to ML: Hypothesis testing with interval data [29], regression based on Choquet’s capacities [66], partial order ranking [11, 28], computation of the Vapnik-Chervonenkis dimension of Choquet integral [34], etc. Beyond, numerous remaining questions are worthy:

- Can we propose an optimizer for the Exploitation-Exploration trade-off (defined in the multi-armed bandit problem [8], as well as in most online learning settings) on the basis of the various imprecision measures available in BFT?
- Is there a mean to reduce the imprecision of a source of information, by means of a convex ℓ^1 -penalized optimization, as proposed in [10] for raw signals?

³ Interestingly enough, this couple of short sentences have raised numerous reviewer’s comments, among which few of them are worthy being discussed here:

- A first comment addressed the fact that in this evolution of ML, the problem formulation remains in the language of statistics, and question its probabilistic interpretation. As a matter of fact, the involvement of a risk function, defined as the expectation of loss function, that is justified to be approximated by the empirical risk on the basis of the assumption that the data are i.i.d., definitely roots machine learning to statistic, multivariate analysis and measure theory; even if traditional hypothesis testing or parametric probabilistic modeling is not involved, and even if optimization methods are used to solve the problem.
- Another comment requires being cited: “*Which ‘solid foundations’? This assertion seems just the repetition of a cliché contrasting the ‘solid foundations’ of probability and statistics to an alleged lack of such foundations for non-probabilistic approaches. We are never told what these ‘solid foundations’ actually are.*” The ‘solid foundations’ of the statistical approaches to ML are described at lenght in Vapnik’s *Statistical learning theory* [69] (and that are sketched in the previous bullet), which in addition to providing a axiomatization of the problem, provides metrics to objectively evaluate experimental results as well as generalization bounds. I have never said that non-probabilistic approaches lack of such foundations, as this anonymous reviewer have misread it; yet I believe these foundations have not been proven to adapt to the current ML problems that are already formalized in the literature.

- As a penalized optimization amounts to finding a trade-off between two criteria (the loss functions and a regularity measure), does it make sense to tackle such type of ML problems in the framework of multicriteria optimization?

6 Conclusion

Finally, even if the modern ML mainstream is less interested in the classification/clustering problems that were a natural application field to BFT, there remain few open questions which would advantageously benefit from BFT. Among these open questions, few are based on well-established interpretations which are affiliated to probability or to set theories (cooperation framework), while some requires the BFT to move forwards, to accept wider interpretations (Choquet's capacities) or to be inserted into wider frameworks (lower probabilities). Among these open questions, few are tentatively addressed by leading researchers who points the directions by providing successful first results (imprecise ranking, computer vision, likelihood interpretation of plausibility), while others remain unaddressed (robust uncertainty principle, similarities between bias-variance and specificity-consistency trade-offs). Finally, despite separating from AI to the profit of applied mathematics, ML still provides an interesting playground to BFT researchers. Yet, most importantly, on few specific ML open questions, we can even expect significant BFT contributions.

References

- [1] http://en.wikipedia.org/wiki/K-means_clustering
- [2] http://en.wikipedia.org/wiki/Empirical_risk_minimization
- [3] <http://www.chalearn.org/>
- [4] Benjamini, Y., Hochberg, Y.: Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J. Royal Stat.Soc. (B)* pp. 289–300 (1995)
- [5] Bezdek, J.C., Ehrlich, R., Full, W.: Fcm: The fuzzy c-means clustering algorithm. *Computers & Geosciences* 10(2), 191–203 (1984)
- [6] Blum, A., Mitchell, T.: Combining labeled and unlabeled data with co-training. In: 11th conference on computational learning theory. pp. 92–100 (1998)
- [7] Boser, B.E., Guyon, I.M., Vapnik, V.N.: A training algorithm for optimal margin classifiers. In: 15th workshop on Computational learning theory. pp. 144–152 (1992)
- [8] Bubeck, S., Cesa-Bianchi, N.: Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends in Machine Learning* 5, 1–122 (2012)
- [9] Burger, T., Aran, O., Caplier, A.: Modeling hesitation and conflict: a belief-based approach for multi-class problems. In: ICMLA. pp. 95–100 (2006)
- [10] Candès, E.J., Romberg, J., Tao, T.: Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *Information Theory, IEEE Transactions on* 52(2), 489–509 (2006)
- [11] Cheng, W., Rademaker, M., De Baets, B., Hüllermeier, E.: Predicting partial orders: Ranking with abstention. In: Machine Learning and Knowledge Discovery in Databases, pp. 215–230. Springer (2010)
- [12] Coifman, R.R., Lafon, S.: Diffusion maps. *Applied and computational harmonic analysis* 21(1), 5–30 (2006)
- [13] Côme, E., Oukhellou, L., Dencœur, T., Aknin, P.: Mixture model estimation with soft labels. In: Soft Methods for Handling Variab. and Imprec., pp. 165–174 (2008)
- [14] Cornuéjols, A., Miclet, L.: What is the place of machine learning between pattern recognition and optimization? (2009)
- [15] Courty, N., Flamary, R., Tuia, D.: Domain adaptation with regularized optimal transport. In: Machine Learning & Knowl. Disc. in Databases, pp. 274–289 (2014)
- [16] Cuzzolin, F.: Belief functions: Theory and applications (2014)
- [17] Cuzzolin, F., Gong, W.: Belief modeling regression for pose estimation. In: 16th International Conference on Information Fusion. pp. 1398–1405 (2013)
- [18] Cuzzolin, F., Gong, W.: A belief-theoretical approach to example-based pose estimation. *IEEE Transactions on Fuzzy Systems* (under revision)
- [19] Dempster, A.: Upper and lower probabilities induced by a multivalued mapping. *Annals of Mathematical Statistics* 38, 325–339 (1967)
- [20] Dempster, A.P.: A generalization of bayesian inference. *Journal of the Royal Statistical Society. Series B (Methodological)* pp. 205–247 (1968)
- [21] Dempster, A.P., Laird, N.M., Rubin, D.B.: Maximum likelihood from incomplete data via the em algorithm. *J. of the Royal Statistical Soc. (B)* pp. 1–38 (1977)

- [22] Denoeux, T.: A k-nearest neighbor classification rule based on dempster-shafer theory. *Systems, Man and Cybernetics, IEEE Transactions on* 25(5), 804–813 (1995)
- [23] Denoeux, T.: A neural network classifier based on dempster-shafer theory. *Systems, Man and Cybernetics (A), IEEE Trans.* 30(2), 131–150 (2000)
- [24] Denœux, T.: Maximum likelihood from evidential data: an extension of the em algorithm. In: *Combining Soft Computing and Statistical Methods in Data Analysis*, pp. 181–188 (2010)
- [25] Denœux, T.: Maximum likelihood estimation from uncertain data in the belief function framework. *Knowl. & Data Engin., IEEE Trans. on* 25(1), 119–130 (2013)
- [26] Denœux, T., Masson, M.H.: Belief Functions: Theory and Applications: Proceedings of the 2nd Int. Conf. on Belief Functions, vol. 164 (2012)
- [27] Destercke, S., Burger, T.: Toward an axiomatic definition of conflict between belief functions. *Systems, Man, & Cyb. (B), IEEE Trans.* 43(2), 585–596 (2013)
- [28] Destercke, S.: A pairwise label ranking method with imprecise scores and partial predictions. In: *Machine Learning and Knowledge Discovery in Databases*, pp. 112–127 (2013)
- [29] Destercke, S., Strauss, O.: Kolmogorov-smirnov test for interval data. In: *Information Processing and Management of Uncertainty in Knowledge-Based Systems*. pp. 416–425. Springer (2014)
- [30] Ding, C., He, X.: K-means clustering via principal component analysis. In: *international conference on Machine learning*. p. 29 (2004)
- [31] Dubois, D., Prade, H.: Possibility theory: qualitative and quantitative aspects. In: *Quantified representation of uncertainty and imprecision*, pp. 169–226 (1998)
- [32] Freund, Y., Schapire, R., Abe, N.: A short introduction to boosting. *Journal-Japanese Society For Artificial Intelligence* 14(771-780), 1612 (1999)
- [33] Grabisch, M.: The application of fuzzy integrals in multicriteria decision making. *European journal of operational research* 89(3), 445–456 (1996)
- [34] Hüllermeier, E., Tehrani, A.F.: On the vc-dimension of the choquet integral. In: *Advances on Computational Intelligence*, pp. 42–50. Springer (2012)
- [35] Johnson, R.A., Wichern, D.W., Education, P.: *Applied multivariate statistical analysis*, vol. 4. Prentice hall Englewood Cliffs, NJ (1992)
- [36] Kendall, D.: Foundations of a theory of random sets. *Stoch. geometry* 3(9) (1974)
- [37] Kessentini, Y., Burger, T., Paquet, T.: A dempster–shafer theory based combination of handwriting recognition systems with multiple rejection strategies. *Pattern Recognition* 48(2), 534–544 (2015)
- [38] Kittler, J., Hatef, M., Duin, R.P., Matas, J.: On combining classifiers. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* 20(3), 226–239 (1998)
- [39] Kohlas, J., Monney, P.A.: A mathematical theory of hints(an approach to the dempster-shafer theory of evidence). *L.N. economics & math. systems* (1995)
- [40] Mairal, J., Bach, F., Ponce, J., Sapiro, G.: Online learning for matrix factorization and sparse coding. *The Journal of Machine Learning Research* 11, 19–60 (2010)
- [41] Mallat, S.: Des mathématiques pour l’analyse de données massives, <http://www.academie-sciences.fr/video/v180214.htm>
- [42] Mitchell, T.M.: Generalization as search. *Artificial Intell.* 18(2), 203–226 (1982)
- [43] Nguyen, H.T.: *An introduction to random sets*. CRC press (2006)
- [44] Pal, N.R., Bezdek, J.C.: On cluster validity for the fuzzy c-means model. *Fuzzy Systems, IEEE Transactions on* 3(3), 370–379 (1995)
- [45] Pichon, F., Destercke, S., Burger, T.: A consistency-specificity trade-off to select source behavior in information fusion. *IEEE trans. on System, Man and Cybernetics - Part B* (accepted in 2014, to appear)
- [46] Powell, G., Marshall, D., Smets, P., Ristic, B., Maskell, S.: Joint tracking and classification of airbourne objects using particle filters and the continuous transferable belief model. In: *9th Intern. Conf. on Information Fusion*. pp. 1–8 (2006)
- [47] Quost, B., Masson, M.H., Denœux, T.: Classifier fusion in the dempster–shafer framework using optimized t-norm based combination rules. *International Journal of Approximate Reasoning* 52(3), 353–374 (2011)
- [48] Reineking, T., Schill, K.: Evidential object recognition based on information gain maximization. In: *Belief Functions: Theory and Applications*, pp. 227–236. Springer (2014)
- [49] Ristic, B., Smets, P.: Target identification using belief functions and implication rules. *Aerospace and Electronic Systems, IEEE Trans.* 41(3), 1097–1103 (2005)
- [50] Ristic, B., Smets, P.: Belief function theory on the continuous space with an application to model based classification. In: *IPMU*. pp. 4–9 (2004)

- [51] Ristic, B., Smets, P.: Kalman filters for tracking and classification and the transferable belief model. In: Int. Conf. on Information Fusion (2004)
- [52] Ristic, B., Smets, P.: Target classification approach based on the belief function theory. Aerospace and Electronic Systems, IEEE Trans. 41(2), 574–583 (2005)
- [53] Scholkopf, B., Smola, A.J.: Learning with kernels: support vector machines, regularization, optimization, and beyond. MIT press (2001)
- [54] Settles, B.: Active learning literature survey. Univ. of Wisconsin, tech. rep. (2010)
- [55] Shafer, G.: A mathematical theory of evidence. Princeton University Press (1976)
- [56] Shafer, G., Shenoy, P.P.: Local computation in hypertrees. Tech. report (1991)
- [57] Shenoy, P.P.: Valuation-based systems for bayesian decision analysis. Operations Research 40(3), 463–484 (1992)
- [58] Shenoy, P.P., Shafer, G.: Propagating belief functions with local computations. IEEE Expert 1(3), 43–52 (1986)
- [59] Shenoy, P.P., Shafer, G.: Axioms for probability and belief-function propagation. In: Uncertainty in Artificial Intelligence 4, pp. 169–198. North-Holland (1990)
- [60] Smets, P., Kennes, R.: The transferable belief model. Artif. Int. 66, 191–243 (1994)
- [61] Smets, P.: The transferable belief model and random sets. International Journal of Intelligent Systems 7(1), 37–46 (1992)
- [62] Smets, P.: Belief functions: the disjunctive rule of combination and the generalized bayesian theorem. International Journal of Approximate Reasoning 9(1), 1–35 (1993)
- [63] Smets, P.: Decision making in the tbm: the necessity of the pignistic transformation. International Journal of Approximate Reasoning 38(2), 133–147 (2005)
- [64] Smets, P., Ristic, B.: Kalman filter and joint tracking and classification based on belief functions in the tbm framework. Information Fusion 8(1), 16–27 (2007)
- [65] Storey, J.D., Tibshirani, R.: Statistical significance for genomewide studies. Proceedings of the National Academy of Sciences 100(16), 9440–9445 (2003)
- [66] Tehrani, A.F., Cheng, W., Dembczyński, K., Hüllermeier, E.: Learning monotone nonlinear models using the choquet integral. Mach. learn. 89(1-2), 183–211 (2012)
- [67] Tibshirani, R.: Regression shrinkage and selection via the lasso. Journal of the Royal Statistical Society. Series B (Methodological) pp. 267–288 (1996)
- [68] Vannoorenberghe, P., Smets, P.: Partially supervised learning by a credal em approach. In: ECSQARU, pp. 956–967 (2005)
- [69] Vapnik, V.N., Vapnik, V.: Statistical learning theory, vol. 2. Wiley New York (1998)
- [70] Von Luxburg, U.: A tutorial on spectral clustering. Statistics and computing 17(4), 395–416 (2007)